

Deducing the Density Hales-Jewett Theorem from an infinitary removal lemma

Tim Austin

Abstract

We offer a new proof of Furstenberg and Katznelson's density version of the Hales-Jewett Theorem:

Theorem. *For any $\delta > 0$ there is some $N_0 \geq 1$ such that whenever $A \subseteq [k]^N$ with $N \geq N_0$ and $|A| \geq \delta k^N$, A contains a **combinatorial line**: that is, for some $I \subseteq [N]$ nonempty and $w_0 \in [k]^{[N] \setminus I}$ we have*

$$A \supseteq \{w : w|_{[N] \setminus I} = w_0, w|_I = \text{const.}\}.$$

Following Furstenberg and Katznelson, we first show that this result is equivalent to a ‘multiple recurrence’ assertion for a class of probability measures enjoying a certain kind of stationarity. However, we then give a quite different proof of this latter assertion through a reduction to an infinitary removal lemma in the spirit of Tao [24] (and also its recent re-interpretation in [3]). This reduction is based on a structural analysis of these stationary laws closely analogous to the classical representation theorems for various partial exchangeable stochastic processes in the sense of Hoover [17], Aldous [1, 2] and Kallenberg [18]. However, the underlying combinatorial arguments used to prove this theorem are rather different from those required to work with exchangeable arrays, and involve crucially an observation that arose during ongoing work by a collaborative team of authors [21] to give a purely finitary proof of the above theorem.

Contents

1	Introduction	2
2	Some background from combinatorics	4
3	Some background from probability	5

4	Strongly stationary laws	7
5	Partially insensitive and sated laws	8
6	The structure of sated laws	12
7	The Density Hales-Jewett Theorem for sated laws	19

1 Introduction

In this note we record a new proof of the following result of Furstenberg and Katznelson:

Theorem 1.1 (Density Hales-Jewett Theorem). *For any $\delta > 0$ there is some $N_0 \geq 1$ such that whenever $A \subseteq [k]^N$ with $N \geq N_0$ and $|A| \geq \delta k^N$, A contains a **combinatorial line**: that is, for some $I \subseteq [N]$ nonempty and $w_0 \in [k]^{[N] \setminus I}$ we have*

$$A \supseteq \{w : w|_{[N] \setminus I} = w_0, w|_I = \text{const.}\}.$$

This is the ‘density’ version of the classical Hales-Jewett Theorem of colouring Ramsey Theory. The Hales-Jewett Theorem is one of the central results of Ramsey Theory, partly because many other results in that area can be deduced from it (see, for example, Chapter 2 of Graham, Rothschild and Spencer [16]). Likewise, the above density variant generalizes many other results in density Ramsey Theory, such as the famous theorem of Szemerédi ([22]) and its multidimensional analog, also proved by Furstenberg and Katznelson ([11]).

Following Furstenberg’s discovery in [9] of an alternative proof of Szemerédi’s Theorem via a conversion to a result in ergodic theory, the use of ergodic-theoretic methods to prove results in density Ramsey Theory has become widespread and powerful (see, for example, the survey of Bergelson [7]), and the above result was one of the furthest-reaching consequences of this program. In this paper we follow Furstenberg and Katznelson as far as their re-interpretation of the above result in terms of stochastic processes, but then we offer a new proof of that version of the result.

In order to state the result about stochastic processes into which Furstenberg and Katznelson convert Theorem 1.1, let us first define $[k]^\omega$ to be the **infinite-dimensional**

combinatorial space over the alphabet k :

$$[k]^\omega := \bigcup_{n \geq 1} [k]^n.$$

Theorem 1.2 (Infinitary Density Hales-Jewett Theorem). *For any $\delta > 0$, if μ is a Borel probability measure on $\{0, 1\}^{[k]^\omega}$ for which*

$$\mu\{\mathbf{x} \in \{0, 1\}^{[k]^\omega} : x_w = 1\} \geq \delta \quad \forall w \in \{0, 1\}^{[k]^\omega}$$

then there are some $N \geq 1$, $I \subseteq [N]$ nonempty and $w_0 \in [k]^{[N] \setminus I}$ such that

$$\mu\{\mathbf{x} \in \{0, 1\}^{[k]^\omega} : x_w = 1 \ \forall w \in [k]^N \text{ s.t. } w|_{[N] \setminus I} = w_0, w|_I = \text{const.}\} > 0.$$

It is shown in Proposition 2.1 of [13] that Theorems 1.1 and 1.2 are equivalent. Here we will assume this first step of Furstenberg and Katznelson and concentrate on proving Theorem 1.2, and will also follow their next step (Lemma 4.2 below) to reduce our study to a class of ‘strongly stationary’ probability measures. Unlike the earlier settings of ergodic Ramsey Theory, this stationarity condition is not readily described by a collection of invertible probability-preserving transformations (in particular, it is instead described in terms of a very large semigroup of highly non-invertible transformations, which cannot easily be made invertible by passing to any simple extended system while respecting the relations of the semigroup). Consequently Furstenberg and Katznelson must next impose a collection of invertible transformations ‘by hand’ that describes only a rather weaker subsemigroup of symmetries, and then bring modifications of their older ergodic-theoretic techniques (see, in particular, [10] and [12]) to bear on these.

Here we avoid the introduction of these transformations, and give an analysis purely in terms of the strong stationarity obtained initially. This has much more in common with many of the basic studies of partially exchangeable arrays of random variables, particularly by Kingman, Hoover, Aldous and Kallenberg: see, for example, the recent book of Kallenberg [19] and the references given there (and also the survey [6], which treats this subject in a very similar formalism to the present paper and also describes the relations of those developments to other combinatorial results in extremal hypergraph theory). Ultimately we reduce the problem to an application of the ‘infinitary hypergraph removal lemma’ of Tao [23] (or, more precisely, of a cut-down corollary of that lemma first used in [3] to give a very similar new proof of the Multidimensional Szemerédi Theorem).

We will prove our main structural result for strongly stationary laws as an assertion that any such law is a ‘factor’ of a law with a particularly simple structure

(similar to the structure of the joint distribution of all the ingredients that are introduced for the representation of an exchangeable array), and then this structure will give the reduction to an infinitary removal lemma. This ‘simple structure’ will be introduced in the definition of ‘sated’ laws in Section 5 below.

We note here that bringing this general program to bear on the task of proving Theorem 1.2 would not have been possible without a crucial insight that recently emerged from an ongoing open collaborative project of Bukh, Gowers, Kalai, McCutcheon, O’Donnell, Solymosi and Tao. Their ultimate goal was a purely finitary, combinatorial proof of Theorem 1.1, and as the present paper neared completion this also seemed to have been realized; these developments can be followed on-line ([21]). The critical observation that we have taken from their work drives our proof of Theorem 6.2 below, but it has been translated into a very different lexicon from the finitary work of [21] and we do not attempt to set up a full dictionary here. I am also grateful to Tim Gowers and Terence Tao for helpful suggestions made about earlier drafts of this paper.

2 Some background from combinatorics

We write $[N]$ to denote the discrete interval $\{1, 2, \dots, N\}$ and $\mathcal{P}S$ to denote the power set of S .

Most of our work will consider probabilities on product spaces indexed by the infinite-dimensional combinatorial space $[k]^\omega$ introduced above. We will denote the concatenation of two finite words $u, v \in [k]^\omega$ by either uv or $u \oplus v$. For any fixed finite n we can define an **n -dimensional subspace** of $[k]^\omega$ to be an injection $\phi : [k]^n \hookrightarrow [k]^\omega$ specified as follows: for some integers $0 = N_0 < N_1 < N_2 < \dots < N_n$, nonempty subsets $I_1 \subseteq [N_1]$, $I_2 \subseteq [N_2] \setminus [N_1]$, \dots , $I_n \subseteq [N_n] \setminus [N_{n-1}]$ and fixed words $w_1 \in [k]^{N_1}$, $w_2 \in [k]^{N_2}$, \dots , $w_n \in [k]^{N_n}$ we let $\phi(v_1 v_2 \dots v_n)$ be the word in $[k]^\omega$ of length N_n such that when $N_i < m \leq N_{i+1}$ we have

$$\phi(v_1 v_2 \dots v_n)_m := \begin{cases} w_m & \text{if } m \in \{N_i + 1, N_i + 2, \dots, N_{i+1}\} \setminus I_{i+1} \\ v_i & \text{if } m \in I_{i+1}. \end{cases}$$

In these terms a combinatorial line is simply a 1-dimensional combinatorial subspace.

Similarly, an **infinite-dimensional subspace** (or often just **subspace**) of $[k]^\omega$ is an injection $\phi : [k]^\omega \hookrightarrow [k]^\omega$ specified by the above rule for some infinite sequence

$0 = N_0 < N_1 < N_2 < \dots$ and nonempty $I_{i+1} \subseteq [N_{i+1}] \setminus [N_i]$. It is clear that the collection of all subspaces of $[k]^\omega$ forms a semigroup under composition.

Finally, let us define **letter-replacement maps**: give $i \in [k]$ and $e \subseteq [k]$, for each $N \geq 1$ we define $r_{e,i}^N : [k]^N \rightarrow [k]^N$ by

$$r_{e,i}^N(w)_m := \begin{cases} i & \text{if } w_m \in e \\ w_m & \text{if } w_m \in [k] \setminus e \end{cases}$$

for $m \leq N$, and let

$$r_{e,i} := \bigcup_{N \geq 1} r_{e,i}^N : [k]^\omega \rightarrow [k]^\omega$$

(so clearly $r_{e,i}$ actually takes values in the subset $([k] \setminus (e \setminus \{i\}))^{[k]^\omega}$).

3 Some background from probability

Throughout this paper (X, Σ) will denote a standard Borel measurable space. We shall write $(X^I, \Sigma^{\otimes I})$ for the usual product measurable structure indexed by a set I and $\mu^{\otimes I}$ for the product of a probability measure μ on (X, Σ) . Given a measurable map $\phi : (X, \Sigma) \rightarrow (Y, \Phi)$ to another standard Borel space, we shall write $\phi_\# \mu$ for the resulting pushforward probability measure on (Y, Φ) . We will generally use π_J to denote any coordinate projection from a product space X^I onto its factor X^J for any $I \supseteq J$, and will shorten $\pi_{\{j\}}$ to π_j .

Most of our interest will be in probability measures on the product spaces $(X^{[k]^\omega}, \Sigma^{\otimes [k]^\omega})$ for various standard Borel spaces (X, Σ) . In this paper we will simply refer to these as **laws**, in view of their interpretation as the joint laws of X -valued stochastic processes indexed by $[k]^\omega$. Let us note here that Theorem 1.2 is clearly equivalent to the following superficially more general result, whose formulation will be more convenient for our proof.

Theorem 3.1. *For any $\delta > 0$, if (X, Σ) is a standard Borel space, μ is a law on $(X^{[k]^\omega}, \Sigma^{\otimes [k]^\omega})$ and $A \in X$ is such that $\mu(\pi_w^{-1}(A)) \geq \delta$ for every $w \in [k]^\omega$ then there are some $m \geq 1$ and a combinatorial line $\ell : [k] \hookrightarrow [k]^m$ such that*

$$\mu\left(\bigcap_{i=1}^m \pi_{\ell(i)}^{-1}(A)\right) > 0.$$

If μ is a law and $\phi : [k]^\omega \hookrightarrow [k]^\omega$ is a subspace, then the projected law $(\pi_{\text{image}(\phi)})_\# \mu$ on $X^{[k]^{\text{image}(\phi)}}$ can be canonically identified with another law on $X^{[k]^\omega}$, simply because ϕ itself gives an identification of $[k]^\omega$ with $\text{image}(\phi)$. In this case we will write $\phi^* \mu$ for this new law on $X^{[k]^\omega}$.

Borrowing some notation from ergodic theory, a **factor** of a law μ on $(X^{[k]^\omega}, \Sigma^{[k]^\omega})$ will be a Borel map $\phi : (X, \Sigma) \rightarrow (Y, \Phi)$ to some other standard Borel space (Y, Φ) . To such a map we can associate its inverse-image σ -subalgebra $\phi^{-1}(\Phi) \leq \Sigma$, and it is standard that in the category of Borel spaces, given a Borel probability measure on μ any σ -subalgebra of Σ agrees with the inverse-image σ -subalgebra of some factor ϕ up to modifying by negligible sets (see, for example, Chapter 2 of Glasner [14]). To such a map ϕ we associate the map $\phi^{[k]^\omega} : [k]^\omega \rightarrow [k]^\omega$ corresponding to the coordinate-wise action of ϕ , and will refer to $(\phi^{[k]^\omega})_\# \mu$ as the associated **factor law** of μ . In the opposite direction, if μ arises from a factor of some ‘larger’ law λ via the factor ϕ then we will refer to λ as an **extension of μ through ϕ** .

An **inverse system of laws** comprises an inverse system of standard Borel spaces

$$\dots \xrightarrow{\psi_{(m+1)}^{(m+2)}} (X_{(m+1)}, \Sigma_{(m+1)}) \xrightarrow{\psi_{(m)}^{(m+1)}} (X_{(m)}, \Sigma_{(m)}) \xrightarrow{\psi_{(m-1)}^{(m)}} \dots \xrightarrow{\psi_{(0)}^{(1)}} (X_{(0)}, \Sigma_{(0)})$$

together with a sequence of laws $\mu_{(m)}$ on $(X_{(m)}^{[k]^\omega}, \Sigma_{(m)}^{[k]^\omega})$ such that $((\psi_{(m)}^{(m+1)})^{[k]^\omega})_\# \mu_{(m+1)} = \mu_{(m)}$ for every m . In this case we will define

$$\psi_{(k)}^{(m)} := \psi_{(k)}^{(k+1)} \circ \psi_{(k+1)}^{(k+2)} \circ \dots \circ \psi_{(m-1)}^{(m)}$$

for $k \leq m$, and will sometimes write instead

$$\dots \xrightarrow{\psi_{(m+1)}^{(m+2)}} (X_{(m+1)}^{[k]^\omega}, \Sigma_{(m+1)}^{\otimes [k]^\omega}, \mu_{(m+1)}) \xrightarrow{\psi_{(m)}^{(m+1)}} (X_{(m)}^{[k]^\omega}, \Sigma_{(m)}^{\otimes [k]^\omega}, \mu_{(m)}) \xrightarrow{\psi_{(m-1)}^{(m)}} \dots$$

as a shorthand to denote this overall situation.

Given an inverse sequence as above, then exactly as in standard ergodic theory (see, for example, Examples 6.3 of Glasner [14]) we can construct an **inverse limit** in the form of a standard Borel space $(X_{(\infty)}, \Sigma_{(\infty)})$, a law $\mu_{(\infty)}$ on $(X_{(\infty)}^{[k]^\omega}, \Sigma_{(\infty)}^{\otimes [k]^\omega})$ and a family of factors $\psi_{(m)} : X_{(\infty)} \rightarrow X_{(m)}$ such that $\psi_{(k)} = \psi_{(k)}^{(m)} \circ \psi_{(m)}$ for all $k < m$ and $((\psi_{(k)})^{[k]^\omega})_\# \mu_{(\infty)} = \mu_{(m)}$ for every m . We will use this construction later in the paper.

Related to the notion of a factor is that of a ‘coupling’: given laws μ and ν on $(X^{[k]^\omega}, \Sigma^{\otimes [k]^\omega})$ and $(Y^{[k]^\omega}, \Phi^{\otimes [k]^\omega})$ respectively, a **coupling** of μ and ν is a law λ

on $((X \times Y)^{[k]^\omega}, (\Sigma \otimes \Phi)^{\otimes [k]^\omega})$ whose coordinate projections onto $(X^{[k]^\omega}, \Sigma^{\otimes [k]^\omega})$ and $(Y^{[k]^\omega}, \Phi^{\otimes [k]^\omega})$ are μ and ν respectively. This definition generalizes to couplings of larger collections of laws in the obvious way. We will also have need for a topology on couplings, set up exactly analogously with the ‘joining topology’ of ergodic theory: quite generally, given a countable collection of standard Borel probability spaces $(X_i, \Sigma_i, \mu_i)_{i \in I}$, the space C of all couplings of the μ_i on the product standard Borel space $(\prod_{i \in I} X_i, \bigotimes_{i \in I} \Sigma_i)$ is endowed with the weakest topology with respect to which all the evaluation maps

$$\lambda \mapsto \int_{\prod_{i \in I} X_i} \prod_{i \in F} f_i \circ \pi_i \, d\lambda$$

for collections $f_i \in L^\infty(\mu_i)$ indexed by finite subsets $F \subseteq I$ are continuous. Just as for joinings of probability-preserving systems (as discussed in Chapter 6 of Glasner [14]), the restriction here to couplings of fixed one-dimensional marginals (rather than arbitrary probability measures on the product space) gives that this is a compact topology on C .

4 Strongly stationary laws

We now introduce the special class of laws that will concern us through most of this paper. These are distinguished by satisfying a kind of ‘self-similarity’ in terms of the structure of the index set $[k]^\omega$.

Definition 4.1 (Strong stationarity). *A law μ on $(X^{[k]^\omega}, \Sigma^{[k]^\omega})$ is **strongly stationary** (s.s.) if $\phi^* \mu = \mu$ for every subspace $\phi : [k]^\omega \hookrightarrow [k]^\omega$.*

This can be thought of as the analog appropriate to the present setting of the exchangeability of a family of random variables (or, equivalently, their joint distribution) under an index-set-permuting action of some countable group: see, for example, Section 2.2 of [6], where this abstract definition is set up before being applied to exchangeable arrays (or ‘exchangeable random hypergraphs’, as they are formulated there).

Indeed, the only real difference between the settings of that paper and this is that here our notion of strong stationarity refers to a semigroup of noninvertible self-maps of the underlying index set, for which it seems difficult to find any ‘invertible model’. Furstenberg and Katznelson meet the same difficulty in their original work, and circumvent it by relying instead only on a weaker symmetry to which

they can associate (using a highly arbitrary selection procedure) a collection of invertible probability-preserving transformations. By contrast, we will find that this noninvertibility is of no consequence for our approach below.

Let us next recall Furstenberg and Katznelson's reduction to the case of s.s. laws, contained in Sections 2.3 and 2.4 of [13].

Lemma 4.2. *If Theorem 3.1 holds for all s.s. laws for every $\delta > 0$ then it holds for all laws for every $\delta > 0$.*

Proof We only sketch the argument, referring the reader to [13] for the details. The point is that by applying the Carlson-Simpson Theorem [8] to arbitrarily fine finite coverings of the finite-dimensional spaces of probability distributions on $\{0, 1\}^{[k]^n}$ for increasingly large n , we obtain a subspace $\psi : [k]^\omega \hookrightarrow [k]^\omega$ and an infinite word $w \in [k]^\mathbb{N}$ such that the restricted laws

$$\psi(w|_{[m]} \oplus \cdot)^* \mu$$

converge to a strongly stationary law as $m \rightarrow \infty$, and since *all* one-dimensional marginals of the input law gave probability at least δ to $\{1\}$, the same is true of the limit. Finally, the condition

$$\mu\{\mathbf{x} \in \{0, 1\}^{[k]^\omega} : x_{\ell(i)} = 1 \ \forall i \leq k\} > 0$$

is also finite-dimensional and open for this topology on the space of finite-dimensional distributions, so if it holds for the limit measure it must also hold somewhere for the original measure. \square

Definition 4.3. *If the law μ is s.s. then in particular all the one-dimensional marginals $(\pi_w)_\# \mu$ for $w \in [k]^\omega$ are the same and all the k -dimensional marginals*

$$(\pi_{\ell(1)}, \pi_{\ell(2)}, \dots, \pi_{\ell(k)})_\# \mu$$

*for ℓ a line in $[k]^\omega$ the same. We will refer to these as the **point-marginal** and **line-marginal** of μ and will often denote them by μ° and μ^{line} respectively.*

5 Partially insensitive and sated laws

Definition 5.1 (Partially insensitive σ -algebras). *For any nonempty $e \subseteq [k]$ and a s.s. law μ on $(X^{[k]^\omega}, \Sigma^{\otimes [k]^\omega})$ the **e -insensitive σ -algebra** is the σ -subalgebra $\Phi_e \leq \Sigma$ defined by*

$$\Phi_e := \{A \in \Sigma : 1_A(x_{\ell(i)}) = 1_A(x_{\ell(j)}) \ \forall i, j \in e \text{ for } \mu\text{-a.e. } (x_w)_w \in X^{[k]^\omega}\}.$$

The e -insensitive σ -algebras for different sets e are together referred to as the **partially invariant σ -algebra**. A measurable function f on X is **e -insensitive** if it is Φ_e -measurable.

The law μ is itself **e -insensitive** if $\Phi_e = \Sigma$, that is if $x_{\ell(i)} = x_{\ell(j)}$ for every $i, j \in e$ for μ -a.e. $(x_w)_w \in X^{[k]^\omega}$.

We now also construct a larger collection of σ -algebras from the above, but first must set up some additional notation. These next σ -algebras will be indexed by **up-sets** in $\binom{[k]}{\geq 2}$: that is, families $\mathcal{I} \subseteq \binom{[k]}{\geq 2}$ such that if $u \in \mathcal{I}$ and $[k] \supseteq v \supseteq u$ then also $v \in \mathcal{I}$. For example, given $e \subseteq [k]$ we write $\langle e \rangle := \{u \in \binom{[k]}{\geq 2} : u \supseteq e\}$ (note the non-standard feature of our notation that $e \in \langle e \rangle$ if and only if $|e| \geq 2$): up-sets of this form are **principal**. We will abbreviate $\langle \{i\} \rangle$ to $\langle i \rangle$.

In general, for any up-set $\mathcal{I} \subseteq \binom{[k]}{\geq 2}$ we let $\Phi_{\mathcal{I}} := \bigvee_{e \in \mathcal{I}} \Phi_e$. It is clear from the above definition that if $e \subseteq e'$ then $\Phi_e \supseteq \Phi_{e'}$, so we have $\Phi_e = \Phi_{\langle e \rangle}$.

It is also immediate from the above definition that for any s.s. law μ , $e \in \binom{[k]}{\geq 2}$ and $i, j \in e$ the σ -subalgebras $\pi_i^{-1}(\Phi_e)$ and $\pi_j^{-1}(\Phi_e)$ of $\Sigma^{\otimes k}$ are equal up to μ^{line} -negligible sets, and so we can make the following definition.

Definition 5.2 (Oblique copies). *For each $e \subseteq [k]$ we refer to the common μ^{line} -completion of the σ -subalgebra $\pi_i^{-1}(\Phi_e)$, $i \in e$, as the **oblique copy** of Φ_e , and denote it by Φ_e^\dagger . More generally we shall refer to σ -algebras formed by repeatedly applying \cap and \vee to such oblique copies as **oblique σ -algebras**.*

Clearly if a law is e -insensitive for some e this amounts to a nontrivial simplification of its structure. In general we will analyze an arbitrary law in terms of its possible couplings to insensitive laws through the following definition.

Definition 5.3 (Sated laws). *For a nonempty up-set $\mathcal{I} \subseteq \binom{[k]}{\geq 2}$ and a s.s. law μ on $(X^{[k]^\omega}, \Sigma^{\otimes [k]^\omega})$ with partially insensitive σ -algebras Φ_e , μ is **\mathcal{I} -sated** if for any s.s. extension $\tilde{\mu}$ of μ the factor $\pi : \tilde{x} \mapsto x$ and the σ -subalgebra $\tilde{\Phi}_{\mathcal{I}}$ are relatively independent under $\tilde{\mu}^\circ$ over the σ -subalgebra $\pi^{-1}(\Phi_{\mathcal{I}})$.*

The law μ is **fully sated** if it is \mathcal{I} -sated for every such \mathcal{I} .

Clearly not all laws are sated, but it turns out that we can recover the advantage of working with a sated law by passing to an extension. The following theorem is closely analogous to a similar ‘satedness’ result to appear in [4], and is also closely related to older results on ‘pleasant’ and ‘isotropized’ extensions of probability-preserving systems in [3, 5].

Theorem 5.4 (Sated extension). *Every s.s. law has a fully sated s.s. extension.*

In light of this it will suffice to prove Theorem 3.1 for fully sated s.s. laws μ . We will finish this section by proving Theorem 5.4, and then in the next section we will derive some useful consequences of full satedness for the structure of μ^{line} before using these to complete the reduction of Theorem 3.1 to an infinitary removal lemma in Section 7.

Lemma 5.5 (Partially sated extension). *For any up-set $\mathcal{I} \subseteq \binom{[k]}{\geq 2}$, every s.s. law μ has an s.s. extension that is \mathcal{I} -sated.*

Proof This proceeds by an infinitary ‘energy increment’ argument: we build a tower of extensions of μ each ‘closer’ to \mathcal{I} -satedness than its predecessor and so that the resulting inverse limit is exactly \mathcal{I} -sated.

Let $(f_r)_{r \geq 1}$ be a countable subset of the L^∞ -unit ball $\{f \in L^\infty(\mu^\circ) : \|f\|_\infty \leq 1\}$ that is dense in this ball for the L^2 -norm, and let $(r_i)_{i \geq 1}$ be a member of $\mathbb{N}^{\mathbb{N}}$ in which every non-negative integer appears infinitely often.

We will now construct an inverse sequence

$$\dots \xrightarrow{\psi_{(m+1)}^{(m+2)}} (X_{(m+1)}^{[k]^\omega}, \Sigma_{(m+1)}^{\otimes [k]^\omega}, \mu_{(m+1)}) \xrightarrow{\psi_{(m)}^{(m+1)}} (X_{(m)}^{[k]^\omega}, \Sigma_{(m)}^{\otimes [k]^\omega}, \mu_{(m)}) \xrightarrow{\psi_{(m-1)}^{(m)}} \dots$$

starting from $(X_{(0)}, \Sigma_{(0)}) = (X, \Sigma)$ and $\mu_{(0)} = \mu$ such that each $(X_{(m+1)}^{[k]^\omega}, \Sigma_{(m+1)}^{\otimes [k]^\omega}, \mu_{(m+1)})$ is obtained by coupling to $(X_{(m)}^{[k]^\omega}, \Sigma_{(m)}^{\otimes [k]^\omega}, \mu_{(m)})$ a new law μ' on some $((X')^{[k]^\omega}, (\Sigma')^{\otimes [k]^\omega})$ such that for this new law we have $\Sigma' = \Phi'_{\mathcal{I}}$.

Suppose that we have already obtained $(X_{(m)}^{[k]^\omega}, \Sigma_{(m)}^{\otimes [k]^\omega}, \mu_{(m)})$ for $0 \leq m \leq m_1$. We consider two separate cases:

- If there is some further extension

$$\pi : (\tilde{X}^{[k]^\omega}, \tilde{\Sigma}^{\otimes [k]^\omega}, \tilde{\mu}) \rightarrow (X_{(m_1)}^{[k]^\omega}, \Sigma_{(m_1)}^{\otimes [k]^\omega}, \mu_{(m_1)})$$

such that

$$\|E_{\tilde{\mu}^\circ}(f_{r_{m_1}} \circ \psi_{(0)}^{(m_1)} \circ \pi \mid \tilde{\Phi}_{\mathcal{I}})\|_2^2 > \|E_{\mu_{(m_1)}^\circ}(f_{r_{m_1}} \circ \psi_{(0)}^{(m_1)} \mid \Phi_{(m_1), \mathcal{I}})\|_2^2 + 2^{-m_1},$$

then choose a particular such extension such that the increase

$$\|E_{\tilde{\mu}^\circ}(f_{r_{m_1}} \circ \psi_{(0)}^{(m_1)} \circ \pi \mid \tilde{\Phi}_{\mathcal{I}})\|_2^2 - \|E_{\mu_{(m_1)}^\circ}(f_{r_{m_1}} \circ \psi_{(0)}^{(m_1)} \mid \Phi_{(m_1), \mathcal{I}})\|_2^2$$

is at least half its supremal possible value over such extensions. Now by restricting to the possibly smaller extension of $(X_{(m_1)}^{[k]^\omega}, \Sigma_{(m_1)}^{\otimes [k]^\omega}, \mu_{(m_1)})$ given by replacing $(\tilde{X}, \tilde{\Sigma})$ with its factor generated by π and the σ -algebra $\tilde{\Phi}_{\mathcal{I}}$, we may assume that $\tilde{\mu}$ is itself obtained as a coupling of $\mu_{(m_1)}$ to a law μ' for which the σ -algebra $\Phi'_{\mathcal{I}}$ is full, and now we let $(X_{(m_1+1)}, \Sigma_{(m_1+1)}) := (\tilde{X}, \tilde{\Sigma})$, $\mu_{(m_1+1)} := \tilde{\mu}$ and $\psi_{(m_1)}^{(m_1+1)} := \pi$.

- If, on the other hand, for every further extension π as above we have

$$\|\mathbb{E}_{\tilde{\mu}^\circ}(f_{r_{m_1}} \circ \psi_{(0)}^{(m_1)} \circ \pi \mid \tilde{\Phi}_{\mathcal{I}})\|_2^2 \leq \|\mathbb{E}_{\mu_{(m_1)}^\circ}(f_{r_{m_1}} \circ \psi_{(0)}^{(m_1)} \mid \Phi_{(m_1), \mathcal{I}})\|_2^2 + 2^{-m_1},$$

then we simply set $\psi_{(m_1)}^{(m_1+1)} := \text{id}_{X_{(m_1)}}$.

Finally, let $(X_{(\infty)}, \Sigma_{(\infty)}, \mu_{(\infty)}^\circ)$ be the inverse limit probability space of

$$\dots \xrightarrow{\psi_{(m+1)}^{(m+2)}} (X_{(m+1)}, \Sigma_{(m+1)}, \mu_{(m+1)}^\circ) \xrightarrow{\psi_{(m)}^{(m+1)}} (X_{(m)}, \Sigma_{(m)}, \mu_{(m)}^\circ) \xrightarrow{\psi_{(m-1)}^{(m)}} \dots,$$

$\mu_{(\infty)}$ the inverse limit of the measures $\mu_{(m)}$ and $\psi_{(m)} : X_{(\infty)} \rightarrow X_{(m)}$ the resulting factor maps. It is clear from the above construction that the whole σ -algebra $\Sigma_{(\infty)}$ is generated up to $\mu_{(\infty)}^\circ$ -negligible sets by $\Phi_{(\infty), \mathcal{I}}$ and $\psi_{(0)}$, since $\Phi_{(\infty), \mathcal{I}}$ contains every $\psi_{(m)}^{-1}(\Phi_{(m), \mathcal{I}})$. To show that $\mu_{(\infty)}$ is \mathcal{I} -sated, let $\tilde{\mu}$ under $\pi : \tilde{X} \rightarrow X_{(\infty)}$ be any further extension of $\mu_{(\infty)}$, and suppose that $f \in L^\infty(\mu_{(\infty)})$. We will complete the proof by showing that

$$\mathbb{E}_{\tilde{\mu}^\circ}(f \circ \pi \mid \tilde{\Phi}_{\mathcal{I}}) = \mathbb{E}_{\mu_{(\infty)}^\circ}(f \mid \Phi_{(\infty), \mathcal{I}}) \circ \pi.$$

By construction, this f may be approximated arbitrarily well in $L^2(\mu_{(\infty)}^\circ)$ by finite sums of the form $\sum_p g_p \cdot h_p$ with g_p being bounded and $\Phi_{(\infty), \mathcal{I}}$ -measurable and h_p being bounded and $\psi_{(0)}$ -measurable, and now by density we may also restrict to using h_p that are each a scalar multiple of some $f_{r_p} \circ \psi_{(0)}$, so by continuity and multilinearity it suffices to prove the above equality for one such product $g \cdot (f_r \circ \psi_{(0)})$. Since g is $\Phi_{(\infty), \mathcal{I}}$ -measurable and $\tilde{\Phi}_{\mathcal{I}} \supseteq \pi^{-1}(\Phi_{(\infty), \mathcal{I}})$, this requirement now reduces to

$$\mathbb{E}_{\tilde{\mu}^\circ}(f_r \circ \psi_{(0)} \circ \pi \mid \tilde{\Phi}_{\mathcal{I}}) = \mathbb{E}_{\mu_{(\infty)}^\circ}(f_r \circ \psi_{(0)} \mid \Phi_{(\infty), \mathcal{I}}) \circ \pi,$$

and this in turn will follow if we only show that

$$\|\mathbb{E}_{\tilde{\mu}^\circ}(f_r \circ \psi_{(0)} \circ \pi \mid \tilde{\Phi}_{\mathcal{I}})\|_2^2 = \|\mathbb{E}_{\mu_{(\infty)}^\circ}(f_r \circ \psi_{(0)} \mid \Phi_{(\infty), \mathcal{I}})\|_2^2.$$

Now, by the martingale convergence theorem we have

$$\|\mathbb{E}_{\mu_{(m)}^\circ}(f_r \circ \psi_{(0)}^{(m)} \mid \Phi_{(m),\mathcal{I}})\|_2^2 \uparrow \|\mathbb{E}_{\mu_{(\infty)}^\circ}(f_r \circ \psi_{(0)} \mid \Phi_{(\infty),\mathcal{I}})\|_2^2$$

as $m \rightarrow \infty$. It follows that if

$$\|\mathbb{E}_{\tilde{\mu}^\circ}(f_r \circ \psi_{(0)} \circ \pi \mid \tilde{\Phi}_{\mathcal{I}})\|_2^2 > \|\mathbb{E}_{\mu_{(\infty)}^\circ}(f_r \circ \psi_{(0)} \mid \Phi_{(\infty),\mathcal{I}})\|_2^2$$

then for some sufficiently large m we would have $r_m = r$ (since each integer appears infinitely often as some r_m) but

$$\begin{aligned} & \|\mathbb{E}_{\mu_{(m+1)}^\circ}(f_{r_m} \circ \psi_{(0)}^{(m+1)} \mid \Phi_{(m+1),\mathcal{I}})\|_2^2 - \|\mathbb{E}_{\mu_{(m)}^\circ}(f_r \circ \psi_{(0)}^{(m)} \mid \Phi_{(m),\mathcal{I}})\|_2^2 \\ & \leq \|\mathbb{E}_{\mu_{(\infty)}^\circ}(f_r \circ \psi_{(0)} \mid \Phi_{(\infty),\mathcal{I}})\|_2^2 - \|\mathbb{E}_{\mu_{(m)}^\circ}(f_r \circ \psi_{(0)}^{(m)} \mid \Phi_{(m),\mathcal{I}})\|_2^2 \\ & < \frac{1}{2} \left(\|\mathbb{E}_{\tilde{\mu}^\circ}(f_r \circ \psi_{(0)}^{(\infty)} \circ \pi \mid \tilde{\Phi}_{\mathcal{I}})\|_2^2 - \|\mathbb{E}_{\mu_{(m)}^\circ}(f \circ \psi_{(0)}^{(m)} \mid \Phi_{(m),\mathcal{I}})\|_2^2 \right) \end{aligned}$$

and also

$$\|\mathbb{E}_{\tilde{\mu}^\circ}(f_r \circ \psi_{(0)}^{(\infty)} \circ \pi \mid \tilde{\Phi}_{\mathcal{I}})\|_2^2 \geq \|\mathbb{E}_{\mu_{(m)}^\circ}(f \circ \psi_{(0)}^{(m)} \mid \Phi_{(m),\mathcal{I}})\|_2^2 + 2^{-m}$$

so contradicting our choice of $\mu_{(m+1)}$ in the first alternative in our construction above. This contradiction shows that we must actually have the equality of L^2 -norms asserted above, as required. \square

Proof of Theorem 5.4 Pick a sequence of up-sets $(\mathcal{I}_m)_{m \geq 1}$ in which each possible up-set appears infinitely often. Now by repeatedly implementing the preceding lemma we can form another tower of extensions

$$\dots \rightarrow (X_{(m+1)}^{[k]^\omega}, \Sigma_{(m+1)}^{\otimes [k]^\omega}, \mu_{(m+1)}) \rightarrow (X_{(m)}^{[k]^\omega}, \Sigma_{(m)}^{\otimes [k]^\omega}, \mu_{(m)}) \rightarrow \dots$$

above $(X^{[k]^\omega}, \Sigma^{\otimes [k]^\omega}, \mu)$ in which every $\mu^{(m)}$ is \mathcal{I}_m -sated. It is now an immediate check that the resulting inverse limit $(\tilde{X}^{[k]^\omega}, \tilde{\Sigma}^{\otimes [k]^\omega}, \tilde{\mu})$ is fully sated. \square

6 The structure of sated laws

Having proved the existence of sated extensions, we will now show how the structure of μ (and particularly of the partially insensitive σ -algebras Φ_e) simplifies for sated systems, before using these results to prove Theorem 3.1 in the next section.

First we need the following lemma.

Lemma 6.1. *If μ is fully sated then for every $i \in e \in \binom{[k]}{\geq 2}$, if $f \in L^\infty(\mu^\circ)$ is e -insensitive then*

$$\mathbb{E}_{\mu^\circ} \left(f \mid \bigvee_{j \in [k] \setminus e} \Phi_{\{i,j\}} \right) = \mathbb{E}_{\mu^\circ} \left(f \mid \bigvee_{j \in [k] \setminus e} \Phi_{e \cup \{j\}} \right).$$

Proof Clearly

$$\mathbb{E}_{\mu^\circ} \left(f \mid \bigvee_{j \in [k] \setminus e} \Phi_{e \cup \{j\}} \right)$$

is always $(\bigvee_{j \in [k] \setminus e} \Phi_{\{i,j\}})$ -measurable. It will therefore suffice to show that if $f \in L^\infty(\mu^\circ)$ is e -insensitive and orthogonal to the σ -algebra $\bigvee_{j \in [k] \setminus e} \Phi_{e \cup \{j\}}$ then it is actually orthogonal to $\bigvee_{j \in [k] \setminus e} \Phi_{\{i,j\}}$. We prove this by contradiction, so suppose for one such f that we could find some bounded functions h_j for $j \in [k] \setminus e$ such that each h_j is $\Phi_{\{i,j\}}$ -measurable and

$$\int_X f \cdot \prod_{j \in [k] \setminus e} h_j \, d\mu^\circ = \kappa \neq 0.$$

Re-writing this inner product condition at the level of the whole law μ it simply reads that

$$\int_{X^{[k]^\omega}} f(x_w) \cdot \prod_{j \in [k] \setminus e} h_j(x_w) \, \mu(d\mathbf{x}) = \kappa$$

for any fixed $w \in [k]^\omega$. However, now we apply first the e -insensitivity of f to deduce that also

$$\begin{aligned} \int_{X^{[k]^\omega}} f(x_w) \cdot \prod_{j \in [k] \setminus e} h_j(x_{r_{e,i}(w)}) \, \mu(d\mathbf{x}) \\ = \int_{X^{[k]^\omega}} f(x_{r_{e,i}(w)}) \cdot \prod_{j \in [k] \setminus e} h_j(x_{r_{e,i}(w)}) \, \mu(d\mathbf{x}) = \kappa \end{aligned}$$

for every word w , and now the $\{i, j\}$ -insensitivity of h_j to deduce that

$$\begin{aligned} \int_{X^{[k]^\omega}} f(x_w) \cdot \prod_{j \in [k] \setminus e} h_j(x_{r_{e,j}(w)}) \, \mu(d\mathbf{x}) \\ = \int_{X^{[k]^\omega}} f(x_w) \cdot \prod_{j \in [k] \setminus e} h_j(x_{r_{e,i}(w)}) \, \mu(d\mathbf{x}) = \kappa. \end{aligned}$$

for every word w .

It follows that if we define the probability measure λ on $(X \times X^{[k] \setminus e})^{[k]^\omega}$ to be the joint law under μ of

$$(x_w)_w \mapsto (x_w, (x_{r_{e,j}(w)})_{j \in [k] \setminus e})_w$$

then all of its coordinate projections onto individual copies of X are still just μ° , the projection

$$\pi : (y_w, (z_{j,w})_{j \in [k] \setminus e})_w \rightarrow (y_w)_w$$

has $\pi_\# \lambda = \mu$ and the projections

$$\pi_j : (y_w, (z_{j,w})_{j \in [k] \setminus e})_w \rightarrow (z_{j,w})_w$$

are λ -almost surely $(e \cup \{j\})$ -insensitive. Therefore through the first coordinate projection π the law λ defines an extension of μ , and the above inequality gives a non-zero inner product for f with some product of $(e \cup \{j\})$ -insensitive functions under λ over $j \in [k] \setminus e$ that we can express as

$$\int_{X^{[k]^\omega}} (f \circ \pi) \cdot \prod_{j \in [k] \setminus e} (h_j \circ \pi_j) d\lambda = \kappa.$$

Now λ may not be stationary, but at least its marginals onto all individual copies of X in $(X \times X^{[k] \setminus e})^{[k]^\omega}$ are equal to μ° . It follows that we can re-run the appeal to the Carlson-Simpson Theorem in Lemma 4.2 to obtain a subspace $\psi : [k]^\omega \hookrightarrow [k]^\omega$ and an infinite word $w \in [k]^\mathbb{N}$ such that the pulled-back measures

$$\psi(w|_{[m]} \oplus \cdot)^* \lambda$$

converge in the coupling topology on $(X \times X^{[k] \setminus e})^{[k]^\omega}$ (recall that for couplings of fixed marginals this is compact) to a s.s. measure $\tilde{\mu}$. Since μ was already strongly stationary, we must still have $\pi_\# \tilde{\mu} = \mu$, and by the definition of the coupling topology as the weakest for which integration of fixed product functions is continuous it follows that we must still have, firstly, that

$$\int_{X^{[k]^\omega}} (f \circ \pi) \cdot \prod_{j \in [k] \setminus e} (h_j \circ \pi_j) d\tilde{\mu} = \kappa,$$

and secondly that the coordinate projections π_j are $(e \cup \{j\})$ -insensitive under $\tilde{\mu}$, since this is equivalent to the assertion that for any $A \in \Sigma$, $i \in e$ and line $\ell : [k] \hookrightarrow [k]^\omega$ we have

$$\int_{(X \times X^{[k] \setminus e})^{[k]^\omega}} 1_A(z_{j, \ell(i)}) \cdot 1_{X \setminus A}(z_{j, \ell(j)}) \tilde{\mu}(d\mathbf{z}) = 0$$

and this is clearly closed in the coupling topology.

Therefore we have found a s.s. extension $\tilde{\mu}$ of μ through some factor map ξ under which

$$E_{\tilde{\mu}^\circ} \left(f \circ \xi \mid \bigvee_{j \in [k] \setminus e} \tilde{\Phi}_{e \cup \{j\}} \right) \neq 0.$$

By satedness, it follows that in fact

$$E_{\mu^\circ} \left(f \mid \bigvee_{j \in [k] \setminus e} \Phi_{e \cup \{j\}} \right) \neq 0,$$

contradicting the condition that f be orthogonal to this σ -algebra. \square

Example The idea behind the above proof may be made clear by an explication of the special case $k = 3$, $i = 2$ and $e = \{1, 2\}$. In this case we wish to prove that if μ is a fully sated s.s. law on $X^{[3]^\omega}$ and $f \in L^\infty(\mu^\circ)$ is $\{1, 2\}$ -insensitive then

$$E_{\mu^\circ}(f \mid \Phi_{\{2,3\}}) = E_{\mu^\circ}(f \mid \Phi_{\{1,2,3\}}),$$

and so we suppose that the right-hand side above is zero and prove that the left-hand side is also zero. Arguing by contradiction, we suppose otherwise and in this case let

$$h := E_{\mu^\circ}(f \mid \Phi_{\{2,3\}}).$$

As a $\Phi_{\{2,3\}}$ -measurable function, this h must be $\{2, 3\}$ -insensitive, and so the condition $h \neq 0$ implies (from the definition of h) that

$$\int_X h^2 d\mu^\circ = \int_X f h d\mu^\circ \neq 0,$$

so we have obtained a nontrivial inner product between the $\{1, 2\}$ -insensitive function f and the $\{2, 3\}$ -insensitive function h . We wish to deduce from this that f actually has a non-zero inner product with some $\{1, 2, 3\}$ -insensitive function. For a general s.s. law μ this does not follow, but using in turn the strong stationarity of μ , the $\{1, 2\}$ -insensitivity of f and then the $\{2, 3\}$ -insensitivity of h we can write

$$\begin{aligned} 0 \neq \int_X f h d\mu^\circ &= \int_{X^{[k]^\omega}} f(x_w) h(x_w) = \int_{X^{[k]^\omega}} f(x_{r_{1,2}(w)}) h(x_{r_{1,2}(w)}) \mu(d\mathbf{x}) \\ &= \int_{X^{[k]^\omega}} f(x_w) h(x_{r_{1,2}(w)}) \mu(d\mathbf{x}) \\ &= \int_{X^{[k]^\omega}} f(x_w) h(x_{r_{3,2}(r_{1,2}(w))}) \mu(d\mathbf{x}) \\ &= \int_{X^{[k]^\omega}} f(x_w) h(x_{222\dots 2}) \mu(d\mathbf{x}) \end{aligned}$$

for any $w \in [k]^\omega$, where $22 \cdots 2$ has the same word-length as w , and this implicitly defines a non-trivial coupling of μ to a process that is indexed by $\{2\}^\omega \subset [3]^\omega$ and which can now be re-interpreted simply as a $\{1, 2, 3\}$ -insensitive law. Applying the Carlson-Simpson Theorem to construct from this a similarly nontrivial coupling that is itself s.s. gives a contradiction with the additional condition that μ be \mathcal{I} -sated for $\mathcal{I} = \{1, 2, 3\}$. \triangleleft

The usefulness of satedness for proving Theorem 3.1 will rest on the following property.

Theorem 6.2. *If $e \subseteq [k]$ is nonempty, μ is fully sated and $f_i \in L^\infty(\mu^\circ)$ for $i \in e$ then*

$$\int_{X^k} \prod_{i \in e} f_i \circ \pi_i \, d\mu^{\text{line}} = \int_{X^k} \prod_{i \in e} \mathbb{E}_{\mu^\circ} \left(f_i \mid \bigvee_{l \in e \setminus \{i\}} \Phi_{\{i, l\}} \right) \circ \pi_i \, d\mu^{\text{line}}.$$

Proof We will prove this by contradiction, assuming that the desired equality fails for some choice of $f_i \in L^\infty(\mu^\circ)$ and constructing from this an extension of μ witnessing that it is not sated. For convenience let us temporarily write $\Xi_i := \bigvee_{l \in e \setminus \{i\}} \Phi_{\{i, l\}}$ (so $\Xi_i = \Phi_{\langle i \rangle}$ when $e = [k]$).

Indeed, given such f_i we can write

$$\begin{aligned} & \int_{X^k} \prod_{i \in e} f_i \circ \pi_i \, d\mu^{\text{line}} - \int_{X^k} \prod_{i \in e} \mathbb{E}_{\mu^\circ}(f_i \mid \Xi_i) \circ \pi_i \, d\mu^{\text{line}} \\ &= \sum_{j \in e} \int_{X^k} \left(\prod_{i \in e, i < j} f_i \circ \pi_i \right) \cdot (f_j \circ \pi_j - \mathbb{E}_{\mu^\circ}(f_j \mid \Xi_j) \circ \pi_j) \cdot \left(\prod_{i \in e, i > j} \mathbb{E}_{\mu^\circ}(f_i \mid \Xi_i) \circ \pi_i \right) \, d\mu^{\text{line}}, \end{aligned}$$

and so if this is nonzero then there is some choice of $j \in e$ for which

$$\int_{X^k} \left(\prod_{i \in e, i < j} f_i \circ \pi_i \right) \cdot (f_j \circ \pi_j - \mathbb{E}_{\mu^\circ}(f_j \mid \Xi_j) \circ \pi_j) \cdot \left(\prod_{i \in e, i > j} \mathbb{E}_{\mu^\circ}(f_i \mid \Xi_i) \circ \pi_i \right) \, d\mu^{\text{line}} \neq 0.$$

Now for each $i \in e \setminus \{j\}$ recall that $r_{j,i} : [k]^\omega \rightarrow [k]^\omega$ is the letter-replacement map defined by

$$(r_{j,i}(w))_m := \begin{cases} i & \text{if } w_m = j \\ w_m & \text{else.} \end{cases}$$

In view of the strong stationarity of μ , the above inequality implies that

$$\int_{X^{[k]^\omega}} \left(\prod_{i \in e, i < j} f_i \circ \pi_{r_{j,i}(w)} \right) \cdot (f_j \circ \pi_w - \mathbb{E}_{\mu^\circ}(f_j \mid \Xi_j) \circ \pi_w) \cdot \left(\prod_{i \in e, i > j} \mathbb{E}_{\mu^\circ}(f_i \mid \Xi_i) \circ \pi_{r_{j,i}(w)} \right) \, d\mu \neq 0$$

for any $w \in [k]^\omega$ such that $w^{-1}\{j\} \neq \emptyset$, since then the points $r_{j,s}(w)$ for $s = 1, 2, \dots, k$ form a combinatorial line.

It follows that if we define the probability measure λ on $(X \times X^{e \setminus \{j\}})^{[k]^\omega}$ to be the joint law under μ of

$$(x_w)_w \mapsto (x_w, (x_{r_{j,i}(w)})_{i \in e, i < j}, (x_{r_{j,i}(w)})_{i \in e, i > j})_w$$

then all of its coordinate projections onto individual copies of X are still just μ° , the projection

$$\pi : (y_w, (z_{i,w})_{i < j}, (z_{i,w})_{i > j})_w \rightarrow (y_w)_w$$

has $\pi_\# \lambda = \mu$ and the projection

$$\pi_{i_0} : (y_w, (z_{i,w})_{i \in e, i < j}, (z_{i,w})_{i \in e, i > j})_w \rightarrow (z_{i_0,w})_w$$

for $i_0 \in e \setminus \{j\}$ is λ -almost surely $\{i_0, j\}$ -insensitive. Therefore through the first coordinate projection π the law λ defines an extension of μ (not necessarily s.s.), and the above inequality gives a fixed non-zero inner product for f with some product of $\{i, j\}$ -insensitive functions under λ over $i \in e \setminus \{j\}$. Arguing exactly as for Lemma 6.1 we obtain the same kind of correlation with some s.s. extension $\tilde{\mu}$ of μ through some factor map ξ , and so in light of the above nonvanishing integral we have

$$\mathbb{E}_{\tilde{\mu}^\circ}(f_j \circ \xi - \mathbb{E}_{\mu^\circ}(f_j \mid \Xi_j) \circ \xi \mid \tilde{\Xi}_j) \neq 0.$$

By satedness, it follows that in fact

$$\mathbb{E}_{\mu^\circ}(f_j - \mathbb{E}_{\mu^\circ}(f_j \mid \Xi_j) \mid \Xi_j) \neq 0,$$

manifestly giving the desired contradiction. \square

Remark Essentially, it is the use of the letter-replacement maps in the above proof that has been brought to the present paper from the online project [21]. This idea was brought to my attention during discussions with Terence Tao, a more active participant in that project. \triangleleft

We can now give our main structural result for sated laws.

Theorem 6.3. *If μ is a fully sated law and $\mathcal{I}, \mathcal{I}' \subseteq \binom{[k]}{\geq 2}$ are two up-sets then the oblique σ -algebras $\Phi_{\mathcal{I}}^\dagger$ and $\Phi_{\mathcal{I}'}^\dagger$ are relatively independent over $\Phi_{\mathcal{I} \cap \mathcal{I}'}^\dagger$ under μ^{line} .*

Remark This result together with Theorem 5.4 amount to our analog for s.s. laws of the representation theorems for partially exchangeable arrays ([18]). \triangleleft

We will deduce Theorem 6.3 result by induction using the following special case.

Lemma 6.4. *If μ is a fully sated law, $\mathcal{I} \subseteq \binom{[k]}{\geq 2}$ is an up-set and e is a member of $\binom{[k]}{\geq 2} \setminus \mathcal{I}$ of maximal size then the oblique σ -algebras Φ_e^\dagger and $\Phi_{\mathcal{I}}^\dagger$ are relatively independent over $\Phi_{\langle e \rangle \cap \mathcal{I}}^\dagger$ under μ^{line} .*

Proof Suppose that $F_1 \in L^\infty(\mu^{\text{line}}|_{\Phi_e^\dagger})$ and $F_2 \in L^\infty(\mu^{\text{line}}|_{\Phi_{\mathcal{I}}^\dagger})$. It will suffice to show that

$$\int_{X^k} F_1 F_2 d\mu^{\text{line}} = \int_{X^k} \mathbb{E}_{\mu^{\text{line}}}(F_1 | \Phi_{\mathcal{I} \cap \langle e \rangle}^\dagger) \cdot F_2 d\mu^{\text{line}}.$$

Pick $i \in e$ and $f_1 \in L^\infty(\mu^\circ|_{\Phi_e})$ such that $F_1 = f_1 \circ \pi_i \mu^{\text{line}}$ -almost surely.

Let $\{a_1, a_2, \dots, a_q\}$ be the antichain of minimal elements in \mathcal{I} ; this clearly generates \mathcal{I} as an up-set. Since $e \notin \mathcal{I}$ we must have $a_s \setminus e \neq \emptyset$ for each $s \leq q$. Pick $i_s \in a_s \setminus e$ arbitrarily for each $s \leq q$, so that $\Phi_{a_s}^\dagger = \pi_{i_s}^{-1}(\Phi_{a_s})$ up to μ^{line} -negligible sets.

Now, since $\Phi_{\mathcal{I}}^\dagger = \bigvee_{s \leq q} \Phi_{a_s}^\dagger$, F_2 may be approximated arbitrarily well in $L^1(\mu^{\text{line}})$ by sums of products of the form $\sum_p \prod_{s \leq q} \phi_{s,p} \circ \pi_{i_s}$ with $\phi_{s,p} \in L^\infty(\mu^\circ|_{\Phi_{a_s}})$, and so by continuity and linearity it suffices to assume that F_2 is an individual such product term. This represents F_2 as a function of coordinates in X^k indexed only by members of $\{i_1, i_2, \dots, i_q\} \subseteq [k] \setminus e$, and now we appeal to Theorem 6.2 to deduce that

$$\begin{aligned} \int_{X^k} F_1 \cdot \prod_{s \leq q} \phi_{s,p} \circ \pi_{i_s} d\mu^{\text{line}} \\ = \int_{X^k} \mathbb{E}_{\mu^\circ}(f_1 | \bigvee_{j \in [k] \setminus e} \Phi_{\{i,j\}}) \circ \pi_i \cdot \prod_{s \leq q} \phi_{s,p} \circ \pi_{i_s} d\mu^{\text{line}}. \end{aligned}$$

However, now Lemma 6.1 and the fact that f_1 is already Φ_e -measurable imply that

$$\mathbb{E}_{\mu^\circ}(f_1 | \bigvee_{j \in [k] \setminus e} \Phi_{\{i,j\}}) = \mathbb{E}_{\mu^\circ}(f_1 | \bigvee_{j \in [k] \setminus e} \Phi_{e \cup \{j\}}),$$

and since each $e \cup \{j\} \in \mathcal{I}$ (by the maximality of e in $\mathcal{P}[k] \setminus \mathcal{I}$), under π_i this conditional expectation must be identified with $\mathbb{E}_{\mu^{\text{line}}}(F_1 | \Phi_{\mathcal{I} \cap \langle e \rangle}^\dagger)$, as required. \square

Proof of Theorem 6.3 We fix \mathcal{I} and prove this by induction on \mathcal{I}' . If $\mathcal{I}' \subseteq \mathcal{I}$ then the result is clear, so now let e be a minimal member of $\mathcal{I}' \setminus \mathcal{I}$ of maximal size, and let $\mathcal{I}'' := \mathcal{I}' \setminus \{e\}$. It will suffice to prove that if $F \in L^\infty(\mu^{\text{line}}|_{\Phi_{\mathcal{I}'}^\dagger})$ then

$$\mathbb{E}_{\mu^{\text{line}}}(F | \Phi_{\mathcal{I}}^\dagger) = \mathbb{E}_{\mu^{\text{line}}}(F | \Phi_{\mathcal{I} \cap \mathcal{I}'}^\dagger),$$

and furthermore, by approximation, to do so only for F that are of the form $F_1 \cdot F_2$ with $F_1 \in L^\infty(\mu^{\text{line}}|_{\Phi_{\langle e \rangle}^\dagger})$ and $F_2 \in L^\infty(\mu^{\text{line}}|_{\Phi_{\mathcal{I}''}^\dagger})$. However, for these we can write

$$\begin{aligned} \mathbb{E}_{\mu^{\text{line}}}(F | \Phi_{\mathcal{I}}^\dagger) &= \mathbb{E}_{\mu^{\text{line}}}(\mathbb{E}_{\mu^{\text{line}}}(F | \Phi_{\mathcal{I} \cup \mathcal{I}''}^\dagger) | \Phi_{\mathcal{I}}^\dagger) \\ &= \mathbb{E}_{\mu^{\text{line}}}(\mathbb{E}_{\mu^{\text{line}}}(F_1 | \Phi_{\mathcal{I} \cup \mathcal{I}''}^\dagger) \cdot F_2 | \Phi_{\mathcal{I}}^\dagger), \end{aligned}$$

and by the preceding lemma

$$\mathbb{E}_{\mu^{\text{line}}}(F_1 | \Phi_{\mathcal{I} \cup \mathcal{I}''}^\dagger) = \mathbb{E}_{\mu^{\text{line}}}(F_1 | \Phi_{(\mathcal{I} \cup \mathcal{I}'') \cap \langle e \rangle}^\dagger).$$

On the other hand $(\mathcal{I} \cup \mathcal{I}'') \cap \langle e \rangle \subseteq \mathcal{I}''$ (because \mathcal{I}'' contains every subset of $[k]$ that strictly includes e , since \mathcal{I}' is an up-set), and so the preceding lemma promises similarly that

$$\mathbb{E}_{\mu^{\text{line}}}(F_1 | \Phi_{(\mathcal{I} \cup \mathcal{I}'') \cap \langle e \rangle}^\dagger) = \mathbb{E}_{\mu^{\text{line}}}(F_1 | \Phi_{\mathcal{I}''}^\dagger).$$

Therefore the above expression for $\mathbb{E}_{\mu^{\text{line}}}(F | \Phi_{\mathcal{I}}^\dagger)$ simplifies to

$$\begin{aligned} \mathbb{E}_{\mu^{\text{line}}}(\mathbb{E}_{\mu^{\text{line}}}(F_1 | \Phi_{\mathcal{I}''}^\dagger) \cdot F_2 | \Phi_{\mathcal{I}}^\dagger) &= \mathbb{E}_{\mu^{\text{line}}}(\mathbb{E}_{\mu^{\text{line}}}(F_1 \cdot F_2 | \Phi_{\mathcal{I}''}^\dagger) | \Phi_{\mathcal{I}}^\dagger) \\ &= \mathbb{E}_{\mu^{\text{line}}}(\mathbb{E}_{\mu^{\text{line}}}(F | \Phi_{\mathcal{I}''}^\dagger) | \Phi_{\mathcal{I}}^\dagger) = \mathbb{E}_{\mu^{\text{line}}}(F | \Phi_{\mathcal{I} \cap \mathcal{I}''}^\dagger) = \mathbb{E}_{\mu^{\text{line}}}(F | \Phi_{\mathcal{I} \cap \mathcal{I}'}^\dagger), \end{aligned}$$

by the inductive hypothesis applied to \mathcal{I}'' and \mathcal{I} , as required. \square

7 The Density Hales-Jewett Theorem for sated laws

Since it is clear that the assertion of Theorem 3.1 holds for a s.s. law if it holds for any extensions of that law, by Theorem 5.4 it suffices to prove Theorem 3.1 in case μ is fully sated.

In this case Theorem 6.2 and Theorem 6.3 together give quite a detailed picture of the joint distribution of the factors $\Phi_{\mathcal{I}}^\dagger$ under μ^{line} , and it turns out that this structure is enough to enable a proof of that theorem along the same lines as for the multidimensional Szemerédi Theorem in [3]. In particular, Theorem 3.1 now follows from an ‘infinitary removal lemma’ essentially identical to that used in [3] (Proposition 6.1 of that paper), which was in turn based on Tao’s ‘infinitary hypergraph removal lemma’ in [24], with some modifications to fit the context of a proof of multiple recurrence. The version we will use below is lifted almost verbatim from [3], and is amenable to an identical proof from Theorem 6.3 as for that result from Corollary 5.2 of [3], so we only state the result here.

Proposition 7.1 (Infinitary removal lemma). *Suppose that μ^{line} is the line marginal of a s.s. law μ , and so has the structure described by Theorem 6.3, and that $\mathcal{I}_{i,j}$ for $i = 1, 2, \dots, d$ and $j = 1, 2, \dots, k_i$ are collections of up-sets in $\binom{[k]}{\geq 2}$ such that $[k] \in \mathcal{I}_{i,j} \subseteq \langle i \rangle$ for each i, j , and suppose further that the sets $A_{i,j} \in \Phi_{\mathcal{I}_{i,j}}$ are such that*

$$\mu^{\text{line}}\left(\prod_{i=1}^d \left(\bigcap_{j=1}^{k_i} A_{i,j}\right)\right) = 0.$$

Then we must also have

$$\mu^\circ\left(\bigcap_{i=1}^d \bigcap_{j=1}^{k_i} A_{i,j}\right) = 0.$$

□

This is proved by an induction on a suitable ordering of the possible collections of up-sets $(\mathcal{I}_{i,j})_{i,j}$, appealing to a handful of different possible cases at different steps of the induction, closely related to the induction on edge-size that underlies the proof of the simplex removal lemma from the finitary hypergraph regularity lemma (see, for example, Gowers [15] or Nagle, Rödl and Schacht [20]). This inductive proof is the reason for the above statement in terms of arbitrary collections of up-sets, but we will need only the special case $k_i = 1$, $\mathcal{I}_{i,1} := \langle i \rangle$ for the proof of Theorem 3.1.

Proof of Theorem 3.1 from Proposition 7.1 As remarked above it suffices to prove Theorem 3.1 for a sated law μ . Given such a law, suppose that $A \in \Sigma$ is such that $\mu^{\text{line}}(A^k) = 0$. Then by Theorem 6.2 we have

$$\mu^{\text{line}}(A^k) = \int_{X^k} \prod_{i=1}^k \mathbb{E}_{\mu^\circ}(1_A | \Phi_{\langle i \rangle}) \circ \pi_i \, d\mu^{\text{line}} = 0.$$

Now the level set $B_i := \{\mathbb{E}_{\mu^\circ}(1_A | \Phi_{\langle i \rangle}) > 0\}$ lies in $\Phi_{\langle i \rangle}$, and the above equality certainly implies that also $\mu^{\text{line}}(B_1 \times B_2 \times \dots \times B_d) = 0$. Now, on the one hand, setting $k_i = 1$, $\mathcal{I}_{i,1} := \langle i \rangle$ and $A_{i,1} := B_i$ for each $i \leq d$, Proposition 7.1 tells us that $\mu(B_1 \cap B_2 \cap \dots \cap B_d) = 0$, while on the other we must have $\mu(A \setminus B_i) = 0$ for each i , and so overall $\mu(A) \leq \mu(B_1 \cap B_2 \cap \dots \cap B_d) + \sum_{i=1}^d \mu(A \setminus B_i) = 0$, as required. □

References

- [1] D. J. Aldous. On exchangeability and conditional independence. In *Exchangeability in probability and statistics (Rome, 1981)*, pages 165–170. North-Holland, Amsterdam, 1982.
- [2] D. J. Aldous. Exchangeability and related topics. In *École d’été de probabilités de Saint-Flour, XIII—1983*, volume 1117 of *Lecture Notes in Math.*, pages 1–198. Springer, Berlin, 1985.
- [3] T. Austin. Deducing the multidimensional Szemerédi Theorem from an infinitary removal lemma. To appear, *J. d’Analyse Math.*
- [4] T. Austin. On characteristic quasifactors subject to some algebraic constraints, and applications. In preparation.
- [5] T. Austin. On the norm convergence of nonconventional ergodic averages. To appear, *Ergodic Theory Dynam. Systems*.
- [6] T. Austin. On exchangeable random variables and the statistics of large graphs and hypergraphs. *Probability Surveys*, (5):80–145, 2008.
- [7] V. Bergelson. Ergodic Ramsey Theory – an Update. In M. Pollicott and K. Schmidt, editors, *Ergodic Theory of \mathbb{Z}^d -actions: Proceedings of the Warwick Symposium 1993-4*, pages 1–61. Cambridge University Press, Cambridge, 1996.
- [8] T. J. Carlson. Some unifying principles in Ramsey theory. *Discrete Math.*, 68:117–169, 1988.
- [9] H. Furstenberg. Ergodic behaviour of diagonal measures and a theorem of Szemerédi on arithmetic progressions. *J. d’Analyse Math.*, 31:204–256, 1977.
- [10] H. Furstenberg. *Recurrence in Ergodic Theory and Combinatorial Number Theory*. Princeton University Press, Princeton, 1981.
- [11] H. Furstenberg and Y. Katznelson. An ergodic Szemerédi Theorem for commuting transformations. *J. d’Analyse Math.*, 34:275–291, 1978.
- [12] H. Furstenberg and Y. Katznelson. An ergodic Szemerédi theorem for IP-systems and combinatorial theory. *J. d’Analyse Math.*, 45:117–168, 1985.
- [13] H. Furstenberg and Y. Katznelson. A Density Version of the Hales-Jewett Theorem. *J. d’Analyse Math.*, 57:64–119, 1991.

- [14] E. Glasner. *Ergodic Theory via Joinings*. American Mathematical Society, Providence, 2003.
- [15] W. T. Gowers. Quasirandomness, counting and regularity for 3-uniform hypergraphs. *Combin. Probab. Comput.*, 15(1-2):143–184, 2006.
- [16] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey Theory*. John Wiley & Sons, New York.
- [17] D. N. Hoover. Relations on probability spaces and arrays of random variables. 1979.
- [18] O. Kallenberg. Symmetries on random arrays and set-indexed processes. *J. Theoret. Probab.*, 5(4):727–765, 1992.
- [19] O. Kallenberg. *Probabilistic symmetries and invariance principles*. Probability and its Applications (New York). Springer, New York, 2005.
- [20] B. Nagle, V. Rödl, and M. Schacht. The counting lemma for regular k -uniform hypergraphs. *Random Structures and Algorithms*, to appear.
- [21] Polymath1. A combinatorial approach to Density Hales-Jewett. Online project, viewable at <http://gowers.wordpress.com/>, 2009.
- [22] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith.*, 27:199–245, 1975.
- [23] T. Tao. A quantitative ergodic theory proof of Szemerédi’s theorem. *Electron. J. Combin.*, 13(1):Research Paper 99, 49 pp. (electronic), 2006.
- [24] T. Tao. A correspondence principle between (hyper)graph theory and probability theory, and the (hyper)graph removal lemma. *J. d’Analyse Math.*, 103:1–45, 2007.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES CA 90095-1555, USA

Email: timaustin@math.ucla.edu

URL: <http://www.math.ucla.edu/~timaustin>